Cross-lingual Abstract Meaning Representation Parsing

Marco Damonte and **Shay B. Cohen**
School of Informatics
University of Edinburgh
m.damonte@sms.ed.ac.uk, scohen@inf.ed.ac.uk

arXiv:1704.04539v1 [cs.CL] 14 Apr 2017

Abstract

Abstract Meaning Representation (AMR) annotation efforts have mostly focused on English. In order to train parsers on other languages, we propose a method based on annotation projection, which involves exploiting annotations in a source language and a parallel corpus of the source language and a target language. Using English as the source language, we show promising results for Italian, Spanish, German and Chinese as target languages. Besides evaluating the target parsers on non-gold datasets, we further propose an evaluation method that exploits the English gold annotations and does not require access to gold annotations for the target languages. This is achieved by inverting the projection process: a new English parser is learned from the target language parser and evaluated on the existing English gold standard.

1 Introduction

Abstract Meaning Representation (AMR) parsing converts natural language sentences into their corresponding AMRs (Banarescu et al., 2013). An AMR is a graph with nodes representing the main events and entities mentioned in the sentence and edges representing the semantic relationships between them. Most datasets for Natural Language Processing (NLP) exist only for the English language and AMR parsing is no exception: the only available AMR datasets large enough to train statistical models are mappings between English sentences and AMR graphs.

Cross-lingual techniques can alleviate this problem by exploiting the English gold annotations and an English parser. One way to do this is by annotation projection, where the existing annotations are projected to a target language through a parallel corpus (Hwa et al., 2005).

The stability of AMR across languages and the extent to which it can be considered as an interlingua have been the subject of preliminary discussions. The original AMR definition states that AMR is not an interlingua (Banarescu et al., 2013). Bojar (2014) categorizes the different kinds of divergences in the annotation between English AMRs and Czech AMRs. Xue et al. (2014) showed that structurally aligning Czech and English AMRs is problematic but reached different conclusions for the language pair English-Chinese, showing that their AMRs can be aligned, suggesting that, with tailored annotation guidelines, it could be possible to make AMR behave as an interlingua.

Even though the cross-lingual stability of AMR is still an open question, in this work we assume that AMR is sufficiently stable and enforce this by making the same AMR graph annotated for English valid for the target language as well. Using annotation projection we can then automatically generate large AMR datasets and train parsers for Italian, Spanish, German and Chinese.

Evaluation is another important issue, as we assume that, for new languages, gold-standard data is not available even for a small test set. We therefore learn a new English parser using the target parser's output, projected back to English, as training data. We can now evaluate this parser the gold standard, which is available for English. We speculate that the score we obtain can be used as a proxy for evaluating the target parser. However, we also evaluate the target parser on non-gold ("silver") data, obtained through the same projection process we use for the training data.

Experiments show promising results for all languages tested, with the Italian, German and Span-

ish parsers performing very similarly between each other.

Contributions This work reports the first results for multilingual AMR parsing. The method proposed allows quick prototyping of multilingual AMR parsers, assuming that an NLP pipeline (POS tagging, NER tagging, dependency parsing) is available for the target language. This prototyping does not require manual annotation. Moreover, using the same cross-lingual method, we propose a novel method to evaluate parsers when gold annotations are missing.

2 AMR Parsing

AMR is a meaning representation originally aimed at English sentences. It relies on Propbank (Kingsbury and Palmer, 2002) to define the main predicates in the sentence. AMRs are rooted and directed graphs $G = (V, E)$ with V and E the set of nodes and edges, respectively. Predicates, entities, and concepts mentioned in the sentence are represented as graph nodes and the relationships between them are represented as graph edges. AMR nodes are either Propbank frames or English words, which means that parsers can in some cases use the English tokens in the input sentences, or their stems, to label nodes. A key property of AMRs is that they allow reentrancies, meaning that a node can have multiple incoming edges. This allows AMR to represent coreferences and control verbs. Nodes in the AMR graph are not annotated with the span of words that triggered them, which makes it necessary to use automatic aligners: in the rest of this paper, we call this type of alignments AMR alignments to distinguish them from the word alignments between words in a parallel corpus. Figure 1 shows an example AMR graph, with the dotted lines representing the AMR alignments. In order to evaluate AMR parsers, the Smatch score (Cai and Knight, 2013) is traditionally used. Smatch computes the overlap in terms of recall, precision, and F-score between two graphs by finding the variable alignments between the graphs that maximize the overlap.

The AMR parser we use for the experiments in this paper was recently introduced by Damonte et al. (2017). It runs in linear time in the length of the sentence and allows quick adaptation to new languages. The parser works in a similar fashion to the *arc-eager* dependency parser (Nivre, 2004): the sentence is stored in a buffer, which is scanned once left-to-right and the words in the sentences are either ignored or trigger a subgraph in a stack. Items in the stack can be then connected by edges in such a way that, when all words have been scanned, a single AMR graph is created. The Smatch score on the LDC2015E86 dataset is 0.64, with the state-of-the-art parser scoring 0.67 on the same dataset (Wang et al., 2016).

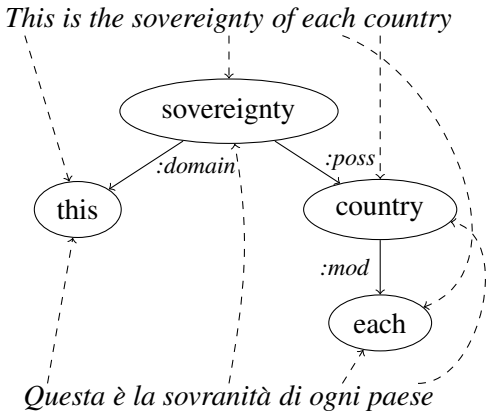

Figure 1: AMR alignments for a English sentence and its Italian translation

3 Cross-lingual Learning for AMR

3.1 AMR Annotation Projection

Projecting the AMR annotation to the target language is straightforward: since AMR is not grounded on the input sentences, we can maintain the AMR annotation as it is. We think of English labels for the graph nodes as using an independent language, which incidentally looks very similar to English. AMR parsing for foreign languages presents therefore an additional complexity: the conversion from tokens to node labels includes an implicit lexical translation.

We also need to project the AMR alignments, necessary to train AMR parsers, as we do not have an AMR aligner for the target languages: this is the part that depends on the actual word tokens in the sentence and we therefore need to use word alignments, as in other annotation projection work, to project the alignments to the target language. An example of AMR graph shared across English and Italian is shown in Figure 1.

Our approach depends on an underlying assumption that we make. Let $A_e(\cdot)$ be the AMR alignment mapping word tokens in the source lan-

guage sentence to the set of AMR nodes that are triggered by it; $A_f(\cdot)$ be the same function for the target language sentence; e_i be a word in the source language; f_j be a word in the target language; v be a node in the AMR graph; and finally, $W(\cdot)$ be an alignment that maps a word in the source language sentence to a subset of words in the target language sentence. The AMR projection assumption is then:

$$\forall i, j, v \quad f_j \in W(e_i) \wedge v \in A_e(e_i) \Rightarrow v \in A_f(f_j) \quad (1)$$

This means that if a source word is word-aligned to a target word and it is AMR aligned with a node in the graph, then the target word is also aligned to that node. In the example of Figure 1, *Questa* is word aligned with *This* and therefore AMR-aligned with the node *this*, and the same logic applies to the other aligned words. The words *is*, *the* and *of* do not generate any AMR nodes, so we ignore their word alignments.

We can use this method to project existing AMR annotations to other languages and then train AMR parsers for them. Unfortunately, there are no parallel corpora with gold AMR annotations available. We therefore use an available English AMR parser to produce these annotations, which we call silver annotations. These are used to train the AMR parser in the target language.

3.2 Evaluation on gold annotations

Once a parser has been learned for a target language (target parser), we are interested in evaluating its performance. We can generate a silver test set in the same way we generated the training set and use this for evaluation. However, the silver test set is influenced by the errors made by the English AMR parser as well as the errors made during the projection, which we later analyze. In order to perform the evaluation on a gold test set, we invert the projection process and train a new English parser from the target parser. We can now evaluate the resulting English parser on a gold test set. Its Smatch score can be then used as a proxy to the score of the target parser.

3.3 Method pipeline

Our method to train and evaluate a parser in a target language is summarized as follows:

1. Parse the English side of a parallel corpus with an English AMR parser.
2. Project the parsed annotation to the target side of the parallel corpus (Equation 1).
3. Train a parser for the target language and evaluate it on silver data.
4. Invert the above procedure to train a *new* English parser from the parser learned in step 3.
5. Evaluate the new English parser on gold data.

Since we rely on several automatic tools, there are several sources of noise in the method: 1) the parsers are trained on silver data obtained by an automatic parser for English ($e \rightarrow f$) or even by an automatic parser trained on data obtained by a parser which in turn was learned with silver data ($f \rightarrow e$); 2) the projection uses noisy word alignments (Padó and Lapata, 2009); 3) the AMR alignments on the source side are also noisy; 4) translation divergences exist between the languages, making it sometimes not possible to project the annotation without loss of information.

4 Experimental setup

We experiment with several languages: Italian, Spanish, German and Chinese. For Italian, Spanish and German we use Europarl (Koehn, 2005), containing around 1.9M sentences for each language pair, and the TED talks corpus (Cettolo et al., 2012) for Chinese, containing around 120K sentences. For the purpose of training the AMR parsers, we extract from each parallel corpus 20,000 sentences for training, 2,000 for development and 2,000 for testing; we collect two such datasets for each language, in order to have non-overlapping datasets for the two stages of the process ($e \rightarrow f$ and $f \rightarrow e$). We use the remainder of the parallel corpora to train the word alignment models. The gold AMR dataset used is LDC2015E86, containing 16,833 training sentences, 1,368 development sentences, and 1,371 testing sentences.

Word alignments were generated using fast_align (Dyer et al., 2013), while AMR alignments were generated with JAMR (Flanigan et al., 2014). As an English AMR parser and a starting point to develop the target parsers we used AMREager (Damonte et al., 2017), which is an open-source AMR parser for English that requires only small modifications for re-use on other languages. The parser requires tokenization, POS tagging, NER tagging and dependency

parsing, which for English, German and Chinese are provided by CoreNLP (Manning et al., 2014). Although CoreNLP also supports Spanish, dependency parsing is not provided. We use Freeling (Carreras et al., 2004) instead. Italian is not supported in CoreNLP: we use Tint (Aprosio and Moretti, 2016), which is a CoreNLP-compatible NLP pipeline for Italian.

5 Results

Besides evaluating on the English gold standard as explained in Section 3.2, we also report silver data evaluation for the target parser. We further define an upper bound for each parser, in which English is used also as the target language: the English side of each parallel corpus is used in both directions, so that we filter out the issues of noisy word alignments and translational divergence. It is necessary to have a different upper bound for each parser because the English sides of each parallel corpus are different. Results are shown in Table 1. The gap between the original English parser (which scores 0.64 on the gold dataset) and the upper bounds shown is approximately 5%, sourced entirely in the effect of re-training on parsed (silver) data. Italian, Spanish and German have comparable performance. For these languages, the gap with the original English parser is around 20%, which is comparable to the 15% gap reported for annotation projection based cross-lingual semantic parsing for French under the combinatory categorial grammar (CCG) framework (Evang and Bos, 2016). The gap with the upper bound is around 15%, a fraction of which is due to noisy alignments, amenable to improvements. For Chinese, the gaps with the English parser and the upper bound are 25% and 18%, respectively. The lower performance of Chinese with respect to the other languages is explained by the different nature and size of the TED talks corpus compared to Europarl. Even within the TED talks corpus, the Chinese-English pair has been shown to be especially challenging (Cettolo et al., 2012). We further analyze the Chinese parser in 6.2.

6 Analysis

Figure 2 shows examples of output parses for all languages tested, including the AMR alignments by-product of the parsing process, that we use to discuss the mistakes made by the parsers.

In the Italian example, the only evident error is that *Infine* (*Lastly*) should be ignored. In the Spanish example, the word *medida* (*measure*) is wrongly ignored: it should be used to generate a child of the node *impact-01*. Some of the *:ARG* roles are also not correct. In the German example, *meines* (*my*) should reflect the fact that the speaker is talking about his own country. Finally, in the Chinese examples, *product* and *increase*, both central to the meaning of the sentence, are ignored. We note that the rest of the AMR is sound, resulting in the correct AMR graph for a different sentence: *Finland exports high technology*.

Most mistakes involve concept identification and in particular relevant words that are erroneously ignored by the parser. This is directly related to the problem of noisy word alignments: the parser learns what words are likely to trigger a node (or a set of nodes) in the AMR by looking at their AMR alignments (which in our approach are induced by the word alignments): if a word is not aligned to any AMR node, it is ignored. Therefore, if an important word is consistently not aligned, the parser will erroneously learn to discard it. This shows that, in order to achieve better parsing results for these languages, more work must be done to allow more accurate alignment projections, which will result in reducing the gap with the upper bound scores.

6.1 Translational divergence

Translational divergence is a known issue with cross-lingual methods. In this section, we look at this type of divergence and discuss how it affects parsing, following the classification used in previous work (Dorr et al., 2002; Dorr, 1994), which identify classes of divergences for several languages. Sulem et al. (2015) also follow the same categorization for French.

To investigate the effects of these differences, we parsed several sentences demonstrating these phenomena to analyze how the AMR parsers coped with them, reported in Figure 3. We do not show the edge labels, as we intend to focus on the identification of the relations, rather than on their labeling. The quality of the AMRs, which is not optimal, is not the main issue here: we only want to assess how the parsers dealt with the different kind of translational divergences.

Categorical. It happens when two languages use different POS tags to express the same meaning. For example, the English sentence *I am jeal-*

	Silver $e \rightarrow f$	Gold $e \rightarrow f \rightarrow e$	Gold $e \rightarrow e \rightarrow e$
Italian	0.45	0.45	0.59
Spanish	0.44	0.44	0.59
German	0.45	0.43	0.59
Chinese	0.36	0.40	0.58

Table 1: Smatch scores for each target parser. "Silver $e \rightarrow f$" is the evaluation of the target parser on data parsed by the existing English parser and projected to the target language, "Gold $e \rightarrow f \rightarrow e$" is the evaluation of the new English parser (learned through the target language parser) on the gold AMR dataset, "Gold $e \rightarrow e \rightarrow e$" is a control experiment where the English side of the corpus is used for both directions. The original English parser scores 0.64 on the gold dataset.

ous of you is translated in Spanish as *Tengo envidia de ti* (*I have jealousy of you*). The English adjective *jealous* is translated in the Spanish noun *envidia*. In Figure 3a we note that the categorical divergence does not create problems, since the parsers correctly recognized that *envidia* (*jealousy*/*envy*) should be used as the predicate, regardless of its POS.

Conflational. It happens when verbs expressed in a language with a single word can be expressed with more words in another language. Two subtypes are distinguished: *manner* and *light verb*. Manner is when a manner verb is mapped to a motion verb plus a manner-bearing word. For example, *We will answer* is translated in the Italian sentence *Noi daremo una riposta* (*We will give an answer*), where *to answer* is translated as *daremo una risposta* (*to give an answer*). Figure 3b shows that the Italian parser generates a sensible output for this sentence by creating a single node labeled *answer-01* for the expression *dare una riposta*. In a light verb conflational divergence, a verb is mapped to a light verb plus an additional meaning unit, such as when *I fear* is translated as *Io ho paura* (*I have fear*) in Italian: *to fear* is mapped to the light verb *ho* (*have*) plus the noun *paura* (*fear*). Figure 3e shows that also this divergence is dealt properly by the Italian parser: *ho paura* correctly triggers the root *fear-01*.

Structural. This type of divergence happens when verb arguments result in different syntactic configurations, for example, due to an additional PP attachment, such as when translating *He entered the house* with *Lui è entrato nella casa* (*He entered in the house*), where the Italian translation has an additional *in* preposition. Also this parse, in Figure 3c, is structurally correct, apart from missing the node *he*.

Head swapping. It occurs when the direction of the dependency between two words is inverted. For example, *I like eating*, where *like* is head of *eating*, becomes *Ich esse gern* (*I eat likingly*) in German, where the dependency is inverted. Unlike the other examples, in this case, the German parser does not cope well with the divergence. Indeed, it is not able to recognize *like-01* as the main concept in the sentence, as shown in Figure 3d.

Thematic. Finally, the parse of Figure 3f has to deal with a thematic divergence, which happens when the semantic roles of a predicate are inverted. In the sentence *I like grapes*, translated to Spanish as *Me gustan uvas*, *I* is the subject in English while *Me* is the object in Spanish. Even though we note an erroneous reentrant edge between *grape* and *I*, the Spanish parser correctly recognizes the *:ARG0* relationship between *like-01* and *I* and the *:ARG1* relationship between *like-01* and *grape*, dealing with the thematic divergence as desired. In this case, the edge labels are important, as this type of divergence is concerned with the semantic roles assigned.

6.2 Analysis of the Chinese Parser

As mentioned, the performance of the Chinese parser is lower than all other parsers. Table 2 shows the differences between the Spanish and the Chinese parsers (silver evaluation), using the submetrics outlined by Damonte et al. (2017) to better investigate where the Chinese parser falls behind. These metrics assess specific problems the AMR parsers need to face such as concept identification, semantic role labeling, and negation detection. We note large gaps for wikification (identification of Wikipedia identifiers for named entities), which

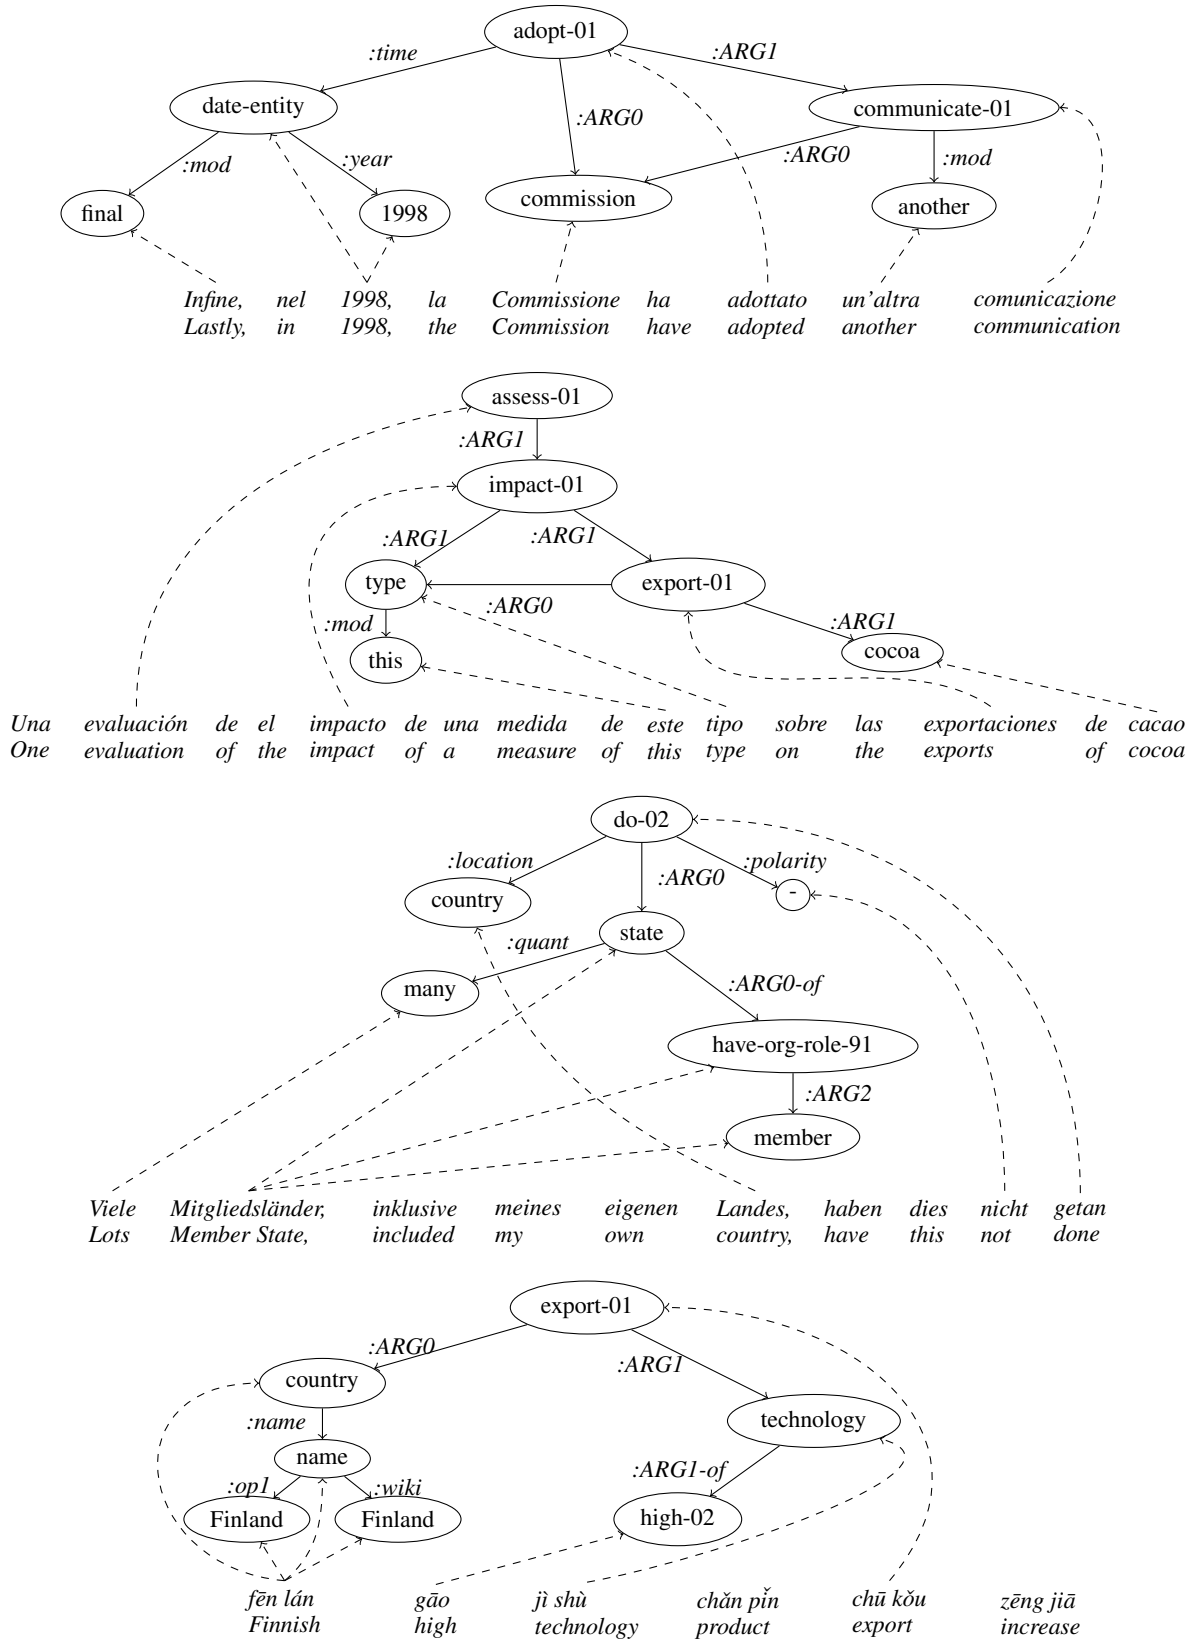

Figure 2: AMR graph and alignments for the Italian translation of *Lastly, in 1998, the Commission adopted a further communication*, the Spanish translation of *A study into the impact of such a measure on the export of cocoa*, the German translation of *Many Member States, including my own, did not do as required* and the Chinese translation of *Finnish exports of high tech products rise*.

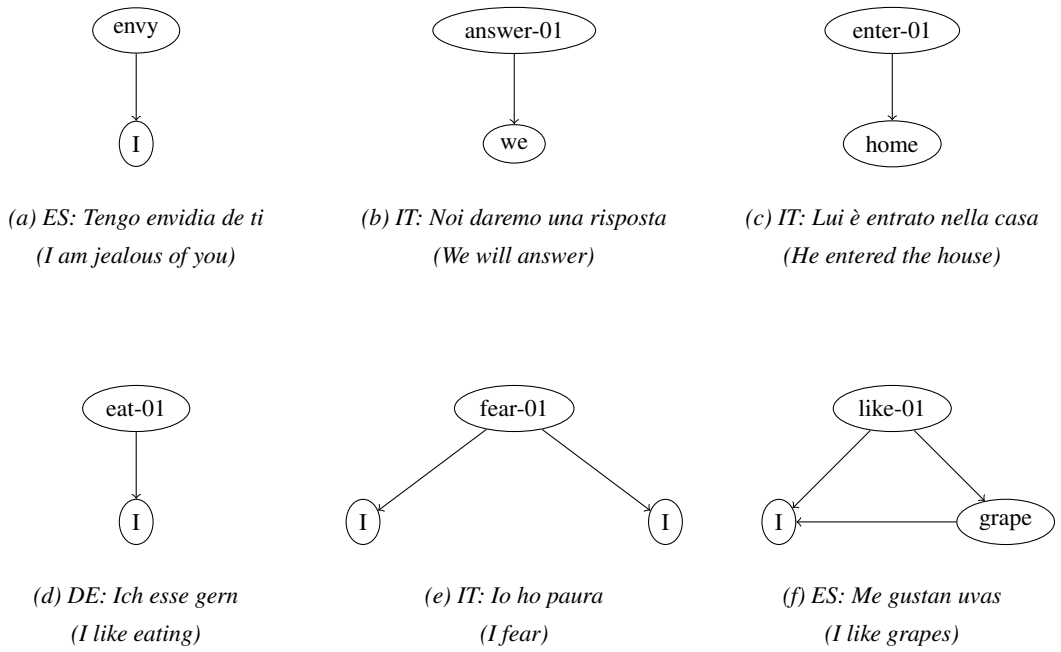

(a) ES: Tengo envidia de ti
(I am jealous of you)

(b) IT: Noi daremo una risposta
(We will answer)

(c) IT: Lui è entrato nella casa
(He entered the house)

(d) DE: Ich esse gern
(I like eating)

(e) IT: Io ho paura
(I fear)

(f) ES: Me gustan uvas
(I like grapes)

Figure 3: Parsing examples in several languages involving common translational divergence phenomena: (a) contains a categorical divergence, (b) and (e) conflational divergences, (c) a structural divergence, (d) an head swapping and (f) a thematic divergence.

Metric	Spanish	Chinese
Smatch	0.44	0.36
Unlabeled	0.51	0.44
No WSD	0.44	0.36
Reentrancy	0.28	0.22
Concepts	0.58	0.45
Named Ent.	0.61	0.63
Wikification	0.58	0.40
Negations	0.22	0.21
SRL	0.38	0.34

Table 2: Detailed comparison between Spanish parser and Chinese parser on silver evaluation. All metrics report F-scores.

however has a small impact on the overall score, and concept identification. We focus on concept identification as its the key step to allow effective multilingual parsing: once the concepts with the correct labels have been identified, the creation of the correct relations can be seen as language independent.

The problem of concept identification and its relationship with noisy alignments was previously discussed. It is present in all languages but more noticeably for Chinese. In the Chinese example of Figure 2, two critical words were ignored in the parsed AMR graph, whereas in the other languages only less important words were erroneously discarded. Table 3 shows, for each parser, the percentage of words seen in training that the parser learned to consider as non-content bearing. Such words are ignored during parsing, resulting in AMRs that lack important concepts. Compared to the original English parser all languages have a high percentage of these words, but the difference is more marked for Chinese. Table 4 highlights specific Chinese words that are treated as non-content bearing when they should not, unlike their English translations, which instead correctly trigger AMR concepts in the English parser.

Parser	%
English	9.4
Italian	28.9
Spanish	28.3
German	25.5
Chinese	41.9

Table 3: Percentage of words seen in training that the parser learned to consider as non-content bearing.

7 Related Work

AMR parsing for languages other than English had made only a few steps forward. In previous

Chinese	English	AMR Concept
xiū jiǎn	trim	trim-02
gēng huàn	replace	replace-01
zhàn	stop	stop-01
biàn yìng	stiffen	stiffen-01
dào lù	way	way
tí mù	topic	topic

Table 4: Examples of Chinese words that are treated as non-content bearing in the Chinese parser but which English translations are treated as content-bearing in the English parser.

work on AMR for other languages (Li et al., 2016; Xue et al., 2014; Bojar, 2014) nodes of the target graph were labeled with either English words or with words in the target language. We instead use the same AMR annotation used for English to the target language, without translating any word. To the best of our knowledge, the only previous work that attempts to automatically parse AMR graphs for non-English sentences is by Vanderwende et al. (2015). Sentences in several languages (French, German, Spanish and Japanese) are parsed into a logical representation, which is then converted to AMR using a small set of rules. A comparison with this work is difficult, as the authors did not report results for the parsers (due to the lack of an annotated corpora) or released their code.

Besides AMR, other semantic parsing frameworks for non-English languages have been investigated (Hoffman, 1992; Cinková et al., 2009; Gesmundo et al., 2009; Evang and Bos, 2016). Evang and Bos (2016) is the most closely related to our work as it uses a projection mechanism similar to ours for CCG. A crucial difference is that, in order to project CCG parse trees to the target languages, they only make use of literal translation, which we argue is not as necessary in our case, since AMR is expected to abstract away from the different syntactic realizations.

Previous work has also focused on assessing the stability across languages of semantic frameworks such as AMR (Xue et al., 2014; Bojar, 2014), UCCA (Sulem et al., 2015) and Propbank (Van der Plas et al., 2010). This work assumes that AMR is sufficiently stable and enforces it by making the same AMR graph annotated for English valid for the target language as well. Supporting evidence comes from the preliminary work on Chinese (Xue et al., 2014) and investigation of the cross-lingual stability of Propbank, on which AMR is partially based (Van der Plas et al., 2010).

Cross-lingual techniques can cope with the lack of labeled data on languages when this data is available in at least one language, normally English. The annotation projection method, which we follow in this work, is only one way to address this problem. It was introduced for dependency parsing (Hwa et al., 2005) but it has also been used for role labeling (Padó and Lapata, 2009) and semantic parsing (Evang and Bos, 2016). Another common thread of cross-lingual work is of the model transfer type, where parameters are shared across languages (Zeman and Resnik, 2008; Cohen et al., 2011; Cohen and Smith, 2009; McDonald et al., 2011; Søgaard, 2011).

8 Conclusion

We proposed a method to overcome the lack of non-English AMR datasets and presented the first results for Italian, Spanish, German and Chinese. Automatic and manual evaluations carried out for these languages are promising and raise hope for further development of AMR parsing for languages other than English. We further proposed a novel way to evaluate the target parsers that does not require manual annotations of the target language. This inversion procedure is not limited to AMR parsing and can be used for other problems in NLP. Finally, we identified weaknesses and the sources of noise in the proposed method, which future work could address.

References

Alessio Palmero Aprosio and Giovanni Moretti. 2016. Italy goes to stanford: a collection of corenlp modules for italian. *arXiv preprint arXiv:1609.06204* .

Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract meaning representation for sembanking. *Linguistic Annotation Workshop* .

Zdenka Urešová Jan Hajic Ondrej Bojar. 2014. Comparing czech and english amrs. In *Workshop on Lexical and Grammatical Resources for Language Processing*.

Shu Cai and Kevin Knight. 2013. Smatch: an evaluation metric for semantic feature structures. *Proceedings of ACL* .

Xavier Carreras, Isaac Chao, Llus Padr, and Muntsa Padr. 2004. Freeling: An open-source suite of language analyzers. In *Proceedings of LREC)*.

Mauro Cettolo, Christian Girardi, and Marcello Federico. 2012. Wit^3: Web inventory of transcribed and translated talks. In *Proceedings of the 16th Conference of the European Association for Machine Translation (EAMT)*. Trento, Italy, pages 261–268.

Silvie Cinková, Josef Toman, Jan Hajic, Kristỳna Cermáková, Václav Klimeš, Lucie Mladová, Jana Šindlerová, Kristỳna Tomšu, and Zdenek Zabokrtskỳ. 2009. Tectogrammatical annotation of the wall street. *The Prague Bulletin of Mathematical Linguistics* .

Shay B. Cohen, Dipanjan Das, and Noah A. Smith. 2011. Unsupervised structure prediction with non-parallel multilingual guidance. In *Proceedings of EMNLP*.

Shay B Cohen and Noah A Smith. 2009. Shared logistic normal distributions for soft parameter tying in unsupervised grammar induction. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, pages 74–82.

Marco Damonte, Shay B Cohen, and Giorgio Satta. 2017. An incremental parser for abstract meaning representation. In *Proceedings of EACL*.

Bonnie J Dorr. 1994. Machine translation divergences: A formal description and proposed solution. *Computational Linguistics* 20(4):597–633.

Bonnie J Dorr, Lisa Pearl, Rebecca Hwa, and Nizar Habash. 2002. Improved word-level alignment: Injecting knowledge about mt divergences. Technical report, DTIC Document.

Chris Dyer, Victor Chahuneau, and Noah A Smith. 2013. A simple, fast, and effective reparameterization of ibm model 2. Proceedings of NAACL-HLT.

Kilian Evang and Johan Bos. 2016. Cross-lingual learning of an open-domain semantic parser. In *Proceedings of COLING*.

Jeffrey Flanigan, Sam Thomson, Jaime G Carbonell, Chris Dyer, and Noah A Smith. 2014. A discriminative graph-based parser for the abstract meaning representation. *Proceedings of ACL* .

Andrea Gesmundo, James Henderson, Paola Merlo, and Ivan Titov. 2009. A latent variable model of synchronous syntactic-semantic parsing for multiple languages. In *Proceedings of CoNLL*. Association for Computational Linguistics.

Beryl Hoffman. 1992. A ccg approach to free word order languages. In *Proceedings of ACL*.

Rebecca Hwa, Philip Resnik, Amy Weinberg, Clara Cabezas, and Okan Kolak. 2005. Bootstrapping parsers via syntactic projection across parallel texts. *Natural language engineering* 11(03):311–325.

Paul Kingsbury and Martha Palmer. 2002. From treebank to propbank. *Proceedings of LREC* .

Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *MT summit*. volume 5, pages 79–86.

Bin Li, Yuan Wen, Lijun Bu, Weiguang Qu, and Nianwen Xue. 2016. Annotating the little prince with chinese amrs. *LAW X Workshop* .

Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *ACL System Demonstrations*. http://www.aclweb.org/anthology/P/P14/P14-5010.

Ryan McDonald, Slav Petrov, and Keith Hall. 2011. Multi-source transfer of delexicalized dependency parsers. In *Proceedings of EMNLP*.

Joakim Nivre. 2004. Incrementality in deterministic dependency parsing. *ACL Workshop on Incremental Parsing: Bringing Engineering and Cognition Together* .

Sebastian Padó and Mirella Lapata. 2009. Cross-lingual annotation projection for semantic roles. *Journal of Artificial Intelligence Research* 36(1):307–340.

Anders Søgaard. 2011. Data point selection for cross-language adaptation of dependency parsers. In *Proceedings of ACL-HLT*.

Elior Sulem, Omri Abend, and Ari Rappoport. 2015. Conceptual annotations preserve structure across translations: A french-english case study. In *Workshop on Semantics-Driven Statistical Machine Translation*.

Lonneke Van der Plas, Tanja Samardžić, and Paola Merlo. 2010. Cross-lingual validity of propbank in the manual annotation of french. In *Linguistic Annotation Workshop*.

Lucy Vanderwende, Arul Menezes, and Chris Quirk. 2015. An amr parser for english, french, german, spanish and japanese and a new amr-annotated corpus. In *Proceedings of NAACL-HLT*. pages 26–30.

Chuan Wang, Sameer Pradhan, Nianwen Xue, Xiaoman Pan, and Heng Ji. 2016. Camr at semeval-2016 task 8: An extended transition-based amr parser. *Proceedings of SemEval* pages 1173–1178.

Nianwen Xue, Ondrej Bojar, Jan Hajic, Martha Palmer, Zdenka Uresova, and Xiuhong Zhang. 2014. Not an interlingua, but close: Comparison of english amrs to chinese and czech. In *Proceedings of LREC*.

Daniel Zeman and Philip Resnik. 2008. Cross-language parser adaptation between related languages. In *IJCNLP*.